An Interop Publication

CONNEXIONS

The Interoperability Report

August 1995

Volume 9, No. 8

*ConneXions—
The Interoperability Report tracks current and emerging standards and technologies within the computer and communications industry.*

In this issue:

ConneXions is published monthly by Interop Company, a division of SOFTBANK Exposition and Conference Company, 303 Vintage Park Drive, Foster City, California, 94404–1138, USA.
Phone: +1 (415) 578-6900
Fax: +1 (415) 525-0194
E-mail: connexions@interop.com

Subscription hotline: 1-800-575-5717
or +1 610-565-6864

ISSN 0894–5926

From the Editor

The *Simple Network Management Protocol* (SNMP) has seen widespread adoption throughout the computer industry. Traditionally, SNMP is used strictly for "network management." However, experience shows that there is also a need to manage other aspects of a distributed computing system. The term "systems management" refers to such items as the underlying operating system, the hardware, the applications, the file system, and so on. Our first article this month is by Bobby Krupczak and describes the current state and future prospects for systems management in the Internet Framework.

Most Internet standards have been developed by a number of *Working Groups* of the *Internet Engineering Task Force* (IETF). The IETF consists of a number of *Areas* focusing on specific technologies. Jim Galvin's reports from the Security Area have become an almost regular feature in *ConneXions,* and we bring you another installment this month. In addition, Joyce Reynolds gives us an overview of another IETF area, namely User Services.

When *ConneXions* began publication in 1987, OSI was considered by many to be the next step in computer networking. For a number of years, OSI versus TCP/IP became the topic of heated arguments, or "protocol wars." About a year ago, the US government released POSIT (a revision of GOSIP) which officially removed its mandate to acquire only OSI protocols. Peter Salus asks if this means that OSI is dead and that the protocol wars are over. Readers with differing opinions are encouraged to respond.

Following the dismantling of the NSFNET backbone in April of this year, the Internet now consists of a number of *Network Access Points* (NAPs) where service providers and other networks exchange traffic. Some aspects of this new system have been described in previous articles in *ConneXions.* This month, Jessica Yu describes the Routing Arbiter *Route Server,* an important component of the new system.

Starting with this issue, our subscription service will be handled by Seybold Publications in Media, Pennsylvania. The 800 number remains unchanged, but for callers outside the United States, please note that you should now call +1 610-565-6864 for subscription questions. The new mailing address and fax appear on page 31.

Finally, if you receive this issue as a free sample at NetWorld+Interop 95 in Paris or Atlanta, we'd like to draw your attention to our special conference discount rates. To subscribe or renew, simply complete the enclosed subscription card and drop it in the mail.

Systems Management and the Internet Management Framework

by Bobby Krupczak, Empire Technologies, Inc.

Introduction

The success, popularity, and importance of using the Internet Management Framework [3,13,14] for network management has been unparalleled by any other management framework. While its efforts were originally focused on the monumental task of managing the "network," managing the "systems" *connected to* the network has become increasingly important (Managing the systems has always been important but industry and the IETF focused on network management first mainly because of its more pressing needs.) Indeed, experience shows that without the correct functioning of key systems, the network becomes unusable from a user, or application, perspective. The need for a consistent method for the management of valuable system resources and mission-critical applications has become increasingly urgent as our dependence on distributed computing continues to grow.

Fortunately, the Internet community has made great strides toward specifying a common interface for performing systems management by defining a number of experimental and standards-track *Management Information Base* (MIB) modules [1,2,8,9] of which the Host Resources MIB [10], on which we focus here, is particularly important for managing systems and their resources. However, in order to fully support the Host Resources MIB, developers of the underlying systems (OS, device drivers, and system daemons) must provide some means of accessing information.

We focus on the strengths and weaknesses of the Host Resources MIB in regards to systems management and, based on implementation experience presented here, identify those areas which require more consistent support by underlying systems. In addition, we outline extensions made within the context of our own private-enterprise MIB which expand the scope of the Host MIB (targeted at "generic" hosts) to provide detailed management of UNIX systems. This article is intended for those interested in network and systems management, and especially those developing hardware, operating systems, and systems applications. It is hoped that these experiences will provide sufficient feedback so that in the future these developers may include better support for accessing systems management information.

This article is organized as follows. First we review systems management and SNMP. Then, we present implementation-based feedback on the Host Resources and provide an overview of our UNIX management extensions.

Systems Management and SNMP

In this section we first define *systems* and *network* management and then provide some motivations for their integration. We then review some of the previous work on systems management within the context of the Internet Management Framework.

The OSI network management literature [5, 6, 7] defines *systems* management as the management of the OSI environment (i.e., the components providing OSI services). It does not necessarily include the underlying hardware and systems software supporting the provision of OSI services. While much of the literature has used the terms *network* and *systems* management interchangeably, they have taken on different meanings than that used in the OSI literature. We use the term *network* management to include the management and operation of a communications subnetwork and the services it provides.

Included in this category are the network interface cards, the protocol software, routers, bridges, etc. We use the term *systems* management to denote the management and operation of the systems in which networked components reside. Included in this category are the underlying operating system, the hardware and devices on which the system operates, its applications, file systems, and its system daemons.

The proliferation of scientific and engineering workstations running complex operating system software has created a nightmare for system administrators. Each operating system defines its own set of system administration utilities and tools, thus requiring administrators to learn multiple interfaces and commands. This lack of a consistent means of management results in increased costs due to training, additional administration personnel, and vendor-specific administration software. However, through the integration of systems and network management, administrators can leverage a common, interoperable and standard interface for the administration of systems. Further, their integration can also permit the leveraging of network management software to perform systems management functions. Lastly, this integration can result in a reduced burden on system support staff and, ultimately, lower management costs.

Interest in systems management via SNMP has been rising since the successful deployment of MIB-II (as early as 1991) although not much has been published regarding its use. However, a few MIBs and papers have been published. The overall applicability of the Internet Management Framework to systems management is explored in [11] along with a discussion on a UNIX systems management MIB. The Host Resources MIB [10] (see Figure 1) defines a set of managed objects useful for systems management.

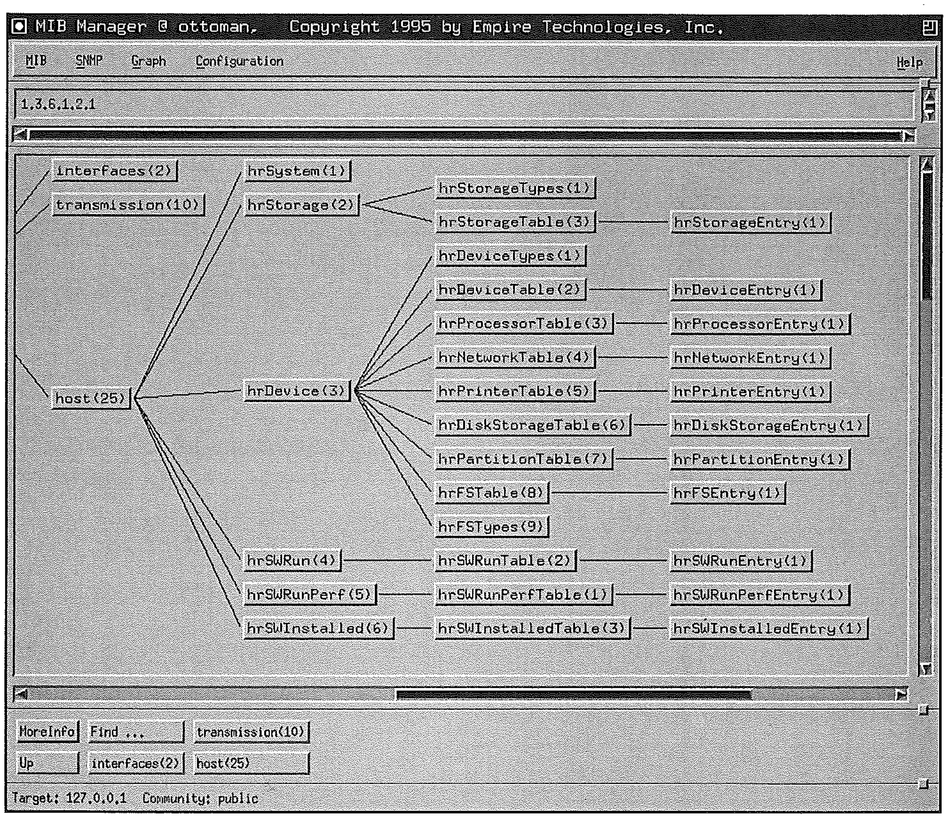

Figure 1: Host Resources MIB, RFC 1514.

continued on next page

Systems Management *(continued)*

It is defined to be operating system independent and contains objects defined for use in the monitoring of devices, storage areas, file systems, and installed software. An excellent overview of this MIB and its potential use can be found in an earlier *ConneXions* article [16]. The monitoring of critical applications has been proposed in several experimental, application-specific MIBs [1, 2, 8, 9, 12]. However, as of yet, no general, flexible framework has been proposed for application management.

Implementation Experiences

In this section we present experiences gained during the implementation of the Host Resources MIB and a private-enterprise UNIX management MIB on two of the major flavors of UNIX (BSD and System V). This implementation has been deployed across a large number of systems in a variety of configurations. We first concentrate on the Host Resources MIB and then discuss its extensions made within our own private-enterprise MIB. Again, while an earlier article ([16]) presents the many benefits of the Host Resources MIB, we mainly focus on the experiences gained during its implementation on real systems.

Device Table

The centerpiece of the Host Resources MIB is the `hrDevice` group and table. This group and table, which provide information about the devices (and their status) contained within a managed system, are extremely useful for asset management. However, complete support for this group is entirely dependent on the underlying operating system and its support for an API to the kernel's information on the hardware and devices currently in operation. While an application can discover some devices by looking through certain operating system configuration files (e.g., `/etc/mnttab` for disks and partitions) or by simply attempting to access known devices, this strategy is unreliable and inherently non-scalable. Accessing static configuration files is often incomplete and prone to errors introduced by configuration changes. Attempting to access known devices is inherently non-scalable since the agent can only report devices it knows about at compile-time. *Some* operating systems *do provide* access to the resources known by the operating system. On these systems, agents can rely on the operating system to inform them of all known devices and their configuration. However, current operating systems do not generally provide runtime access to each device's description, manufacturer, version, etc.—information important for asset management. Some devices, notably SCSI disks and tape drives, do provide such access through well-defined APIs. It is our belief that all devices should provide some well-known API for accessing basic information about its manufacturer, revision, type, etc. and that operating systems should make all relevant device information available to application or systems programmers. On UNIX systems, a system defined *ioctl* would probably be sufficient.

Buffers and the Storage Table

The Host Resources MIB defines a `hrStorage` group and table designed to provide information about the logical storage areas contained within a system. However, the integration of buffer subsystem statistics into the `hrStorageTable` is problematic due to several factors. First, the `hrStorageTable` does not really define a storage type for operating system buffers; a MIB implementor must choose between random-access memory (RAM) or the catch-all "other" for the storage type. Choosing the storage type "RAM" could cause confusion with the system's total RAM, while choosing "other" does not provide much semantic information.

Further, if one chose RAM for the storage-type, there could be overlap if different subsystems define buffers of the same size. For example, in SunOS systems, buffer statistics for the Streams and BSD network subsystems as well as the general I/O buffer cache could all be supported. Clumping all buffer information together as RAM would be confusing and could result in the loss of management information. One solution would be to define additional storage types specifically for buffers.

MIB-II Correlation

The **`hrNetworkTable`** defines a table for network devices contained within the system. One of its columns, **`hrNetworkIfIndex`**, identifies the network device's corresponding MIB-II **`ifTable.ifNumber`** value. The Host Resources MIB *does not* define a return value to use if this network device has no corresponding MIB-II **`ifTable`** entry. For example, many Sun workstations contain an ISDN interface which, technically, is a network interface but many times is not contained in the MIB-II **`ifTable`** because it is not currently used as an IP network interface. In our implementation, we return the value of 0 in such a case. The MIB specification should be modified so that a value of 0 indicates that a network interface is not currently contained in the **`ifTable`**.

System Configuration

The Host Resources MIB defines only a handful of system configuration variables within its **`hrSystem`** group. Included are the initial boot/load device, the number of users, and the maximum number of processes allowed by the system. In our experience, many more configuration variables are possible and needed to properly manage a system (regardless of its operating system). One alternative is to try to incorporate these variables into the **`hrStorage`** table. For example, one could construe file and inode lists (or, for example, their DOS and Windows equivalents) to be storage areas in the system and include their allocation information in the **`hrStorage`** table. However, the same problems arising from the introduction of buffer information in the **`hrStorage`** also arise in this case as well.

Installed Software

The Host Resources MIB defines a table (**`hrSWInstalledTable`**) containing an entry for each software package installed on the system. There are two problems with its design and implementation. First, the table does not contain a column for the software package's version. Instead, the implementor must combine the software version (if it is known) with the package's name. This oversight hinders automated management since human intervention will probably be necessary in order to interpret what portion of the software package's name is the version. Unfortunately, type-casting software versions into SNMP defined types may not be easy since vendors often use incompatible numbering systems.

Second, not all operating systems (e.g., SunOS and other BSD derived systems) support a unified software installation and packaging format. One work-around is to attempt to do a software "discover"; however, it is a time-consuming task and one is then left with the problem of deciding what is a software package and what is not. Does the agent include shell scripts and system binaries as software packages? Microsoft Windows performs a software "discover" by comparing files it finds against a master list. While this strategy might be practical on PCs with limited disk space, it is not scalable to multi-user systems with large amounts of internal disk storage. Further, with such a strategy, the master list of known software is frozen at the time the agent is distributed. In order to provide support for newer software packages, the agent (or its software list) would also need to be updated.

continued on next page

Systems Management *(continued)*

ProductID Variables

The Host Resources MIB defines an object identifier textual convention, `ProductID`, that is intended to identify the manufacturer, model, and version of a hardware item or software package. At present, few hardware or software vendors support it; indeed, conformance with this aspect of the MIB was cause of some debate within the original working group. If vendors begin to comply, it is important that they support the ability to dynamically determine the `ProductID` so as to avoid requiring agents to maintain lists of `ProductID` to vendor mappings. For UNIX systems, a well-known *ioctl* could be defined such that its return value is a device or kernel module's `ProductID`. Application software should support `ProductIDs` as part of installation information.

Disk, Partition, and File System Tables

The Host Resources MIB defines tables containing disk information, disk partitions, and file systems residing within the managed system. Starting with the disk table, managers may then examine that disk's partitions and then examine a given partition's file system. In order to properly support all disks (not just those that have file systems and are mounted), the operating system must provide information about all known devices. Second, it must support manufacturer-independent methods of obtaining configuration information such as capacity, file system layout, etc. Lastly, the `hrFSTable` defines columns identifying the last full and partial backup dates. Unfortunately, there are a myriad (e.g., *bar, tar, dump*) of backup programs and techniques which each store backup times (if at all) in an incompatible manner. What is needed is a consistent method for storing backup times for partitions and file systems. One solution would be to have a partition's backup times stored within the disk's label along with the normal partition and cylinder information.

Traditional Agent Paradigm

The Host Resources MIB continues the tradition of defining MIB objects without endowing the agent or MIB (depending on your perspective) with "intelligence" for distributed self-management. Other MIBs, most notably the *Remote Network Monitoring MIB* [15] and the *Manager-to-Manager MIB* [4], are able to endow agents with intelligence, yet still conform to the philosophies of the Internet Management Framework. The Host Resources MIB, however, defines no traps or self-monitoring capabilities; consequently, managers must (perhaps unavoidably) poll the agent.

Advantages

Despite these issues, the Host Resources MIB is extremely useful for managing systems. Its device and installed-software tables are excellent for asset tracking and management. Its disk, partition, and file system tables coupled with the Host Resources system group are useful for configuration management. Finally, its "standard" nature makes it easier for management station applications to code towards a vendor-independent format.

Host Resources MIB Extensions

Because the Host Resources MIB specification focuses on those aspects common to all computer systems, it may not completely satisfy the management requirements of a particular system (e.g., UNIX). Consequently, we have evolved our own private-enterprise MIB to extend the systems management capabilities of the Host Resources MIB. More specifically, we have geared (at present) our MIB specification to the task of managing UNIX and UNIX-like operating systems. Our extensions have been twofold. First, we have defined and implemented additional MIB objects important to the management of UNIX systems. Second, we have begun defining and implementing "RMON"-like extensions for systems management. We discuss these complimentary directions below.

Our first task was the definition and implementation of a range of MIB objects important to the management of UNIX and UNIX-like operating systems. We added extensive kernel configuration parameters covering everything from file table sizes, swap space, and the BSD and Streams I/O subsystems. Next, we added extensive process information to augment that contained in the Host Resources `hrSWRunTable` and `hrSWRunPerfTable`. For example, we report a process' owner and group as well as its process-ID, size, and nice value. Because of the importance of memory buffers, we created a separate group for buffer statistics and separated them by their respective subsystems (e.g., *strbufs* and *mbufs*). We also added information about a system's valid users and groups as well as information about who is currently logged on and using the system. Since information about valid users and groups can be sensitive (for security purposes), we also added the ability to selectively "turn on or off" support for these groups. (More sophisticated MIB views within the context of the SNMPv2 administration model would be a more general solution though). While the Host Resources MIB tends to focus on configuration and asset management, we added support for performance management by adding MIB objects for tracking kernel, disk, and memory performance statistics. Lastly, we added a range of counters and statistics necessary for monitoring NFS and RPC services due to their importance in distributed environments.

Our second task was the definition of "RMON"-like extensions for systems management whereby we endow agents with intelligence, yet still conform to the goals of the Internet Management Framework. Our motivations are based primarily on the need to provide a scalable systems and network management solution to large, distributed networks of workstations. By endowing an agent with intelligence and autonomy, we can reduce network and system load due to polling and instruct the agent on how to behave in certain circumstances. Our "RMON"-like extensions have taken (so far) two forms.

First, we have defined a monitor table (roughly equivalent to the RMON's `alarmTable`) that enables a manager to instruct an agent to monitor certain MIB objects and to send a TRAP message when some threshold is crossed. Each monitor entry defines a Boolean expression, that when evaluated to True, indicates that an event has occurred. For example, using the monitor table, a manager can instruct an agent to send a TRAP message when a certain disk or file system becomes 95% full; alternatively, a manager can instruct an agent to send a TRAP when an important process (e.g., *sendmail*) dies or becomes a zombie. Second, we are defining and implementing a history mechanism whereby a manager can instruct an agent to sample and store the value of a MIB object over time. Like the RMON MIB [15], we have defined history and history control tables and the notion of *sample buckets*. Sample buckets serve to place an upper bound on the resources the agent devotes to a particular history control entry.

For example, using this facility, managers can instruct the agent to sample and record the length of the run queue, the amount of network I/O, the size of a given process, or the amount of allocated and free buffer space over some interval. This type of functionality (monitoring and recording) is crucial for the proper tuning of UNIX systems. Although we have initially geared our MIB towards the task of managing UNIX, there is nothing in our "RMON"-like extensions that is necessarily specific to UNIX.

continued on next page

Systems Management *(continued)*

Consequently, these extensions may be leveraged for the management of other operating systems. Finally, because these extensions have been patterned after the RMON MIB, we hope to leverage RMON management software.

Conclusion

The Host Resources MIB is a great start toward the goal of integrating network and systems management. What we have focused on is experiences gained from its implementation on two major flavors of UNIX. However, due to its goal of being generic, some potential manageability of UNIX systems is not realized. Consequently, we have evolved our own private-enterprise MIB by adding MIB objects more specific to the task of managing UNIX and by defining "RMON"-like functionality for managing systems. It is hoped that these experiences will guide future system developers in making their systems more "management-ready" for those implementing systems management agents.

References

[1] R. Austein and J. Saperia, "DNS server MIB extensions," RFC 1611, May 1994.

[2] D. Brower, R. Purvy, A. Daniel, M. Sinykin, and J. Smith, "Database management system (RDBMS) management information base (MIB) using SMIv2," RFC 1697, August 1994.

[3] J. D. Case, M. Fedor, M. L. Schoffstall, and C. Davin, "Simple network management protocol," RFC 1157, May 1990.

[4] Jeffrey D. Case, Keith McCloghrie, Marshall T. Rose, and Steven L. Waldbusser, "Manager-to-manager management information base," RFC 1451, April 1993.

[5] Information processing systems: Open systems interconnection basic reference model part 4: Management framework, ISO/IEC 7498-4, October 1989.

[6] Information technology: Open systems interconnection: Systems management — part 1: Object management function, ISO/IEC 10164-1, CCITT X.730, August 1991.

[7] Information technology: Open systems interconnection: Systems management overview, ISO/IEC 10040, CCITT X.701, August 1991.

[8] N. Freed and S. Kille, "Mail monitoring MIB," RFC 1566, January 1994.

[9] N. Freed and S. Kille, "Network services monitoring MIB," RFC 1565, January 1994.

[10] P. Grillo and S. Waldbusser, "Host resources MIB," RFC 1514, September 1993.

[11] B. Krupczak, "Unix systems management via SNMP," In *Proceedings of the IFIP TC6/WG6.6 Third International Symposium on Integrated Network Management,* April 1993.

[12] G. Mansfield and S. Kille, "X.500 directory monitoring MIB," RFC 1567, January 1994.

[13] M. T. Rose, "Concise MIB definitions," RFC 1212, March 1991.

[14] M. T. Rose and K. McCloghrie, "Structure and identification of management information for TCP/IP-based internets," RFC 1155, May 1990.

[15] Steve Waldbusser, "Remote network monitoring management information base," RFC 1271, November 1991.

[16] Steven L. Waldbusser, "A look at the host resources MIB," *ConneXions,* Volume 7, No. 11, November 1993.

[17] *ConneXions,* Two *Special Issues* on Network Management and Network Security, Volume 3, No. 3, March 1989 and Volume 4, No. 8, August 1990.

[18] Case, J. D., Davin, J. R., Fedor, M. S., & Schoffstall, M. L., "Network Management and the Design of SNMP," *ConneXions,* Volume 3, No. 3, March, 1989.

[19] Case, J., McCloghrie, K., Rose, M. T., Waldbusser, S., "The Simple Management Protocol and Framework: Managing the Evolution of SNMP," *ConneXions,* Volume 6, No. 10, October 1992.

[20] Rose, M. T., "Network Management is Simple: You just need the 'Right' Framework." In *Integrated Network Management, II,* Iyengar Krishan and Wolfgang Zimmer, editors, pages 9–25, North Holland, April 1991.

[21] Rose, M. T., *The Simple Book: An Introduction to Internet Management,* Second Edition, Prentice-Hall, ISBN 0-13-177254-6, 1993.

[22] McConnell, J., "DMTF: A Foundation for Systems Management?" *ConneXions,* Volume 8, No. 9, September 1994.

[23] Case, J. D., McCloghrie, K., Rose M. T., & Waldbusser, S. L., "Introduction to version 2 of the Internet-standard Network Management Framework," RFC 1441, April 1993.

[24] Case, J. D., McCloghrie, K., Rose M. T., & Waldbusser, S. L., "Structure of Management Information for version 2 of the Simple Network Management Protocol (SNMPv2)," RFC 1442, April 1993.

[25] Case, J. D., McCloghrie, K., Rose, M. T. & Waldbusser, S. L., "Management Information Base for version 2 of the Simple Network Management Protocol (SNMPv2)," RFC 1450, April 1993.

[26] Case, J. D., McCloghrie, K., Rose, M. T. & Waldbusser, S. L., "Textual Conventions for version 2 of the Simple Network Management Protocol (SNMPv2)," RFC 1443, April 1993.

[27] Case, J. D., McCloghrie, K., Rose, M. T. & Waldbusser, S. L., "Conformance Statements for version 2 of the Simple Network Management Protocol (SNMPv2)," RFC 1444, April 1993.

[28] Galvin, J. M., & McCloghrie, K., "Administrative Model for version 2 of the Simple Network Management Protocol (SNMPv2)," RFC 1445, April 1993.

[29] Rose, M. T., "Network Management: Status and Challenges," *ConneXions,* Volume 7, No. 6, June 1993.

BOBBY KRUPCZAK is Chief Scientist at Empire Technologies, Inc. where he designs and implements intelligent network and systems management agents and managers. He regularly participates on the InteropNet NOC team where has gained valuable insight into the art of network and systems management of large, heterogeneous networks. He can be reached at `rdk@empiretech.com`. For more information see: `http://www.empiretech.com/empiretech/`

Spring IETF Meeting Security Activities

by Jim Galvin, Trusted Information Systems

Introduction

The *Internet Engineering Task Force* (IETF) met in Danvers, Massachusetts, April 3–7, 1995. The IPv6 security requirements received considerable attention during the open *Internet Engineering Steering Group* (IESG) meeting. As reflected in the IP next generation protocol recommendation (RFC 1752), published as a Proposed Standard in January 1995, IPv6 requires authentication and encryption to be available at all times and authentication to be enabled by default.

All of the usual perspectives were expressed, including:

- *Political:* With the export restrictions such as they are, why bother?
- *Technical:* Is DES the best choice for an encryption algorithm?
- *Emotional:* We need security now! We need standards to guarantee interoperability! Do something!

No consensus was reached but none had been expected. It was important, however, for opinions to be expressed, and the IESG should consider them in future deliberations of IP next generation specifications.

Fewer-than-usual security activities occurred at this meeting. The *Privacy Enhanced Mail* (PEM) working group did not meet. Documents are awaiting final review by the group and subsequent submission to the IESG for publication as proposed standards. The PEM and MIME integration protocol being proposed is now called *MIME Object Security Services* (MOSS).

The *Domain Name System Security* (DNSSEC) working group did not meet, although the latest version of the specification is under revision, largely due to the implementation efforts of Trusted Information Systems. The Secure HTTP birds of a feather (BOF) held at the last meeting did not submit a charter (the minimum requirement for becoming a working group). However, it was announced at this meeting that a new working group would be formed, *Web Transport Security,* the scope of which includes the work proposed by the Secure HTTP BOF.

Finally, the IP Security working group met only once during this week, in contrast to its tradition of multiple meetings. As usual, the meeting was lively, and discussions continue on the mailing list. Consensus was reached on requiring the implementation of a hybrid Diffie-Hellman key exchange for the *Internet Key Management Protocol* (IKMP). The Photuris specification will be progressed in the working group as the baseline description of this key exchange. Additional presentations were given on possible modifications to the Photuris exchange and on a framework for the IKMP protocol signaling.

A summary of security-relevant working groups and BOF sessions follows.

Common Authentication Technology (CAT)

The working group charter is being revised to include a scope compatible with *Independent Data Unit Protection* (IDUP) and with future authorization support facilities. The following documents are eligible for the issuance of a working group last call (as indicated) in advance of their submission to the IESG for publication as Proposed Standards.

The *FTP Security* specification has been resurrected with a new editor, March Horowitz. A revised specification is expected to be available soon.

A revision of Version 2 of the GSS specification and C bindings documents is expected to be available soon.

The *Simple Public Key Mechanism* is considered stable at this time, pending resolution of an issue identified regarding X9.44 replay protection, and was placed in working group last call on April 5, prior to its submission to the IESG for publication as a Proposed Standard.

RFC 1510 (Kerberos) is now eligible for advancement to a Draft Standard. One remaining issue relevant to the document's advancement will be resolved on the mailing list.

NetScape gave a presentation on the *Secure Socket Layer.* An Internet-Draft is expected to be available soon.

Discussion continues on the other documents and activities.

GRIP

The *Guidelines and Recommendations for Security Incident Processing* (GRIP) is part of the Operational Requirements Area but is tracked by the Security Area. This new working group has drafted guidelines for security incident response teams. A revised draft is expected to be available soon, at which time a working group last call will be issued. The document will be submitted to the IESG for publication as a Proposed Standard by the next IETF.

Router Requirements

The Router Requirements working group has a document ready to submit to the IESG for publication as a Proposed Standard. It will include a requirement level of SHOULD for strong remote authentication. The use of security by routing protocols is an issue to be resolved before the document advances to Draft Standard.

More information

For more information see: `http://www.ietf.cnri.reston.va.us/`

[Ed.: The next IETF meeting was held in Stockholm, Sweden, July 17–21, 1995. We hope to bring you another IETF security update from the Stockholm meeting in a future issue.]

References

[1] Galvin, J., "Security Awareness Increasing within the IETF," *ConneXions,* Volume 8, No. 9, September 1994.

[2] Galvin, J., "Summer IETF Meeting Security Report," *ConneXions,* Volume 8, No. 10, October 1994.

[3] Galvin, J., "Security Activities at the Winter IETF Meeting," *ConneXions,* Volume 9, No. 4, April 1995.

[4] Galvin, J., "The Deployment of Privacy Enhanced Mail," *ConneXions,* Volume 5, No. 10, October 1991.

[5] Schiller, J., "Issues in Internet Security," *ConneXions,* Volume 7, No. 9, September 1993.

[Ed.: A version of this report appeared in the *Data Security Letter.* For more information contact: `dsl@tis.com`]

JAMES M. GALVIN is a Senior Computer Scientist at Trusted Information Systems (TIS). Dr. Galvin's responsibilities emphasize communications security, especially computer networks, architectures, policies, and procedures. He is a principal in the development of TIS' openly available implementation of Privacy Enhanced Mail. He is very active in the IETF Security Area, serving as Executive Director of the Security Directorate, and past Chair of the OSI Implementor's Workshop Security Special Interest Group, hosted quarterly by the National Institute of Standards and Technology. He received his Ph.D. and M.S. degrees, both in Computer Science, from the University of Delaware in 1988 and 1986, respectively. In 1982, he received his B.S. in Computer Science and Mathematics from Moravian College in Bethlehem, PA. E-mail: `galvin@tis.com`

The User Services Area of the IETF

by Joyce K. Reynolds, ISI

Introduction

When the *Internet Engineering Task Force* (IETF) was first established, it did not immediately create a distinct User Services Area. Since 1991, this area has grown to take its place with other Internet *Engineering Steering Group* (IESG) areas as the importance of the user has increased globally. This area provides an international forum for people interested in all levels of user services, to identify and initiate projects designed to improve the quality of the information available to users of the Internet.

One continuing goal of the User Services Area is to coordinate the development of user information services by clearly and concisely providing documentation information and distribution for the Internet community. FYI (*For Your Information*) RFCs (*Request for Comments*) are introductory and overview documents for network users. Their purpose is to make available general information, rather than the protocol specifications or standards that is typical of other RFCs. FYIs are allied to the RFC series of notes, but provides information about who does what on the Internet. The FYI RFC series has proved a success since its initiation, and its goal is to continue to do so. A current list of FYI RFCs are listed at the end of this article.

Working Groups

The actual projects of the User Services Area are handled by the creation of *Working Groups* (WGs). There are currently ten working groups in this area, as summarized in Table 1 below.

WG	*Chair(s)*	*Mailing List*
HARTS	Scott Stoner Susan Siegfried	`harts@isi.edu`
IDS	Linda Millington Sri Satalur	`ids@merit.edu`
ISN	Jennifer Sellers Jodi Chu	`isn-wg@nasa.gov`
NISI	April Marine	`nisi@merit.edu`
RUN	Sally Hambridge Gary Malkin	`ietf-run@mailbag.intel.com`
SSH	Barbara Fraser	`ssh@cert.org`
TRAINMAT	Jill Foster Mark Prior	`us-wg@nic.near.net`
URI	Larry Masinter	`uri@bunyip.com`
USERGLOS	Gary Malkin	`usergloss@xylogics.com`
USWG	Joyce K. Reynolds	`us-wg@nic.near.net`

Table 1: IETF User Services Area June 1995

• *Humanities and Arts (HARTS):* The HARTS Working Group has identified three goals to be pursued. The first goal is development of a FAQ (*Frequently Asked Questions*) regarding value and role of the arts in the Internet. The second goal is to define tools for artists' organizations on the Internet which will focus on creating, viewing, and storage formats for arts humanities resources. This will include contributions regarding text, sound, still and motion images.

It will address different operating systems, glossary of basic terminology and a bibliography. The third goal is to further define issues surrounding copyright and intellectual property, funding, and other support for arts humanities participation and any other needs identified by the survey.

• *Integrated Directory Services (IDS):* The IDS Working Group is chartered to facilitate the integration and interoperability of current and future directory services into a unified directory service. This work will unite directory services based on a heterogeneous set of directory services protocols (X.500, WHOIS++, etc.). In addition to specifying technical requirements for the integration, the IDS Group will also contribute to the administrative and maintenance issues of directory service offerings by publishing guidelines on directory data integrity, maintenance, security, and privacy and legal issues for users and administrators of directories.

• *Internet School Networking (ISN):* ISN is chartered to address issues related to the connection of primary and secondary schools worldwide to the Internet. The key audiences include network service providers and educational policy makers responsible for network access and use. The key areas of focus for this group are advocacy and articulation.

• *Internet User Glossary (USERGLOS):* USERGLOS is chartered to create an Internet glossary of networking terms and acronyms for the Internet community. The WG will update the existing Internet Users' Glossary (RFC 1392, FYI 18), which is now over two years old. The group will make any necessary corrections and add new terminology which has evolved since the last edition was created.

• *Network Information Services Infrastructure (NISI):* NISI is exploring the requirements for common, shared Internet-wide network information services. The goal is to develop an understanding for what is required to implement an information services "infrastructure" for the Internet.

• *Network Training Materials (TRAINMAT):* TRAINMAT is chartered to enable the research community to make better use of the networked services. Towards this end, the Working Group will work to provide a comprehensive package of "mix and match" training materials for the broad academic community which will: 1) enable user support staff to train users to use the networked services and 2) provide users with self-paced learning material. In the first instance, it will not deal with operational training. This Working Group is the IETF component of a joint TERENA/IETF group working on Network Training Materials.

• *Responsible Use of the Network (RUN):* RUN is chartered to create an etiquette ("netiquette" in network parlance) guide for Internet users. The working group will develop an FYI RFC on responsible use of the Internet and its tools.

• *Site Security Handbook (SSH):* SSH is chartered to create two documents: (1) a revised handbook that will help system and network administrators develop their own site-specific policies and procedures to deal with computer security problems and their prevention and (2) a new handbook for users. The text of these documents will be developed from the existing RFC 1244, plus needed revisions and additions.

continued on next page

The User Services Area of the IETF *(continued)*

- *Universal Resource Identifiers (URI):* URI is chartered to define a set of standards for the encoding of system independent Resource Location and Identification information for the use of Internet information services.

- *User Services (USWG):* The User Services Working Group provides a regular forum for people interested in all user services to identify and initiate projects designed to improve the quality of information available to end-users of the Internet.

References

The FYI RFC Series of Internet Documentation listed below was developed specifically for *Users* (not Wizards!)

[1] FYI 27 "Tools for DNS debugging," (Also RFC 1713), November 1994.

[2] FYI 26 "K–12 Internetworking Guidelines," (Also, RFC 1709), November 1994.

[3] FYI 25 "A Status Report on Networked Information Retrieval: Tools and Groups," (Also, RFC 1689, RTR 13), August 1994.

[4] FYI 24 "How to Use Anonymous FTP," (Also RFC 1635), May 1994.

[5] FYI 23 "Guide to Network Resource Tools," (Also RFC 1580), March 1994.

[6] FYI 22 "FYI on Questions and Answers Answers to Commonly Asked 'Primary and Secondary School Internet User' Questions," (Also RFC 1578), February 1994.

[7] FYI 21 "A Survey of Advanced Usages of X.500," (Also RFC 1491), July 1993.

[8] FYI 20 "FYI on "What is the Internet?," (Also RFC 1462), May 1993.

[9] FYI 19 "FYI on Introducing the Internet—A Short Bibliography of Introductory Internetworking Readings," (Also RFC 1463), May 1993.

[10] FYI 18 "Internet Users' Glossary," (Also RFC 1392), January 1993.

[11] FYI 17 "The Tao of IETF—A Guide for New Attendees of the Internet Engineering Task Force," (Also RFC 1718), November 1994.

[12] FYI 16 "Connecting to the Internet: What Connecting Institutions Should Anticipate," (Also RFC 1359), August 1992.

[13] FYI 15 "Privacy and Accuracy Issues in Network Information Center Databases," (Also RFC 1355), August 1992.

[14] FYI 14 "Technical Overview of Directory Services Using the X.500 Protocol," (Also RFC 1309), March 1992.

[15] FYI 13 "Executive Introduction to Directory Services Using the X.500 Protocol," (Also RFC 1308), March 1992.

[16] FYI 12 "Building a Network Information Services Infrastructure," (Also RFC 1302), February 1992.

[17] FYI 11 "A Catalog of Available X.500 Implementations," (Also RFC 1292), January 1992.

[18] FYI 10 "There's Gold in them thar Networks! or Searching for Treasure in all the Wrong Places," (Also RFC 1402), January 1993.

[19] FYI 9 "Who's Who in the Internet: Biographies of IAB, IESG and IRSG Members," (Also RFC 1336), May 1992.

[20] FYI 8 "Site Security Handbook," (Also RFC 1244), July 1991.

[21] FYI 7 "FYI on Questions and Answers: Answers to Commonly Asked "Experienced Internet User" Questions," (Also RFC 1207), February 1991.

[22] FYI 6 "FYI on the X Window System," (Also RFC 1198), January 1991.

[23] FYI 5 "Choosing a Name for Your Computer," (Also RFC 1178), August 1990.

[24] FYI 4 "FYI on Questions and Answers: Answers to Commonly asked "New Internet User" Questions," (Also RFC 1594), March 1994.

[25] FYI 3 "FYI on Where to Start: A Bibliography of Internetworking Information," (Also RFC 1175), August 1990.

[26] FYI 2 "FYI on a Network Management Tool Catalog: Tools for Monitoring and Debugging TCP/IP Internets and Interconnected Devices," (Also RFC 1470), June 1993.

[27] FYI 1 "F.Y.I. on F.Y.I.: Introduction to the F.Y.I. Notes," (Also RFC 1150), March 1990.

[28] Benford, Steve, "Components of OSI: X.500 Directory Services," *ConneXions,* Volume 3, No. 6, June 1989.

[29] Fälström, Patrik, "The WHOIS++ Directory Service," *ConneXions,* Volume 8, No. 12, December 1994.

[30] Deutsch, P., "The Internet as Service Provider," *ConneXions,* Volume 7, No. 9, September 1993.

[32] Berners-Lee, T., "A Summary of the WorldWideWeb System," *ConneXions,* Volume 6, No. 7, July 1992.

How to get RFCs

Details on obtaining FYI RFCs via FTP or e-mail may be obtained by sending an e-mail message to `rfc-info@isi.edu` with the message body "`help: ways_to_get_rfcs.`" For example:

```
To: rfc-info@isi.edu
Subject: getting rfcs
help: ways_to_get_rfcs
```

JOYCE K. REYNOLDS has been affiliated with USC/Information Sciences Institute since 1979. Ms. Reynolds has contributed to the development of the DARPA Experimental Multimedia Mail System, the Post Office Protocol, the Telnet Protocol, and the Telnet Option Specifications. She helped update the File Transfer Protocol (FTP). Ms. Reynolds is an Internet Engineering Steering Group (IESG) member of the Internet Engineering Task Force (IETF). She established a new informational series of notes for the Internet community: FYI (For Your Information) RFCs. FYI RFCs are documents useful to network users. Their purpose is to make available general and useful information with broad applicability. As chair of the User Services Area Council (USAC), the coordinating committee for the IETF User Services Area, she provides leadership in the international user services arena. Ms. Reynolds received Bachelor of Arts and Master of Arts degrees in the Social Sciences from the University of Southern California (USC). She is the past Associate Editor of the *Internet Society News.* E-mail: `jkrey@isi.edu`.

Protocol Wars: Is OSI Finally Dead?

by Peter H. Salus

If you're "on the net," you know about TCP/IP. But that's not the "official" protocol suite. The *International Organization for Standardization* (ISO) legislated something else (called OSI in 1978) and the National Bureau of Standards (now NIST) went along with it. There's been a lot of strife over the past 17 years.

An end to war

But now, in line with peace in Northern Ireland and in the Middle East, NIST seems to have declared an end to protocol wars after nearly two decades. On September 16, 1994, a notice was entered in the Federal Register for a "new FIPS that renames GOSIP to 'Profiles for Open Systems Internetworking Technologies (POSIT)' FIPS 146-2." This Federal Information Processing Standard will remove the mandate to acquire only OSI protocols.

My guess is that wild cheers have rung out in networking centers all over North America. Publication in the Federal Register marked the beginning of the required 45-day comment period, which (appropriately) ended just after Halloween 1994. Revised publication occurred in May, 1995.

FIPS 146-2 recommends the use of IGOSS-NIST SP 500-217 to agencies wishing to acquire computer networking products based on the OSI standards. For IPS guidance, the FIPS 146-2 makes general reference to the IETF voluntary standards and specific reference to RFC 1610. Finally, the reference to the *Government Network Management Profile* (GNMP) FIPS 179-1 will be removed.

GOSIP

GOSIP required US agencies to procure equipment that could run the ISO–OSI (International Organization for Standardization Open Systems Interconnection) network protocols. As the government is such a large customer, this meant that vendors added OSI (or at least software that looked as though OSI protocols could be added) to their networking software to qualify for agency sales.

GOSIP never required anyone to use OSI, merely to obtain it. Vendors also supplied the TCP/IP suite, and that is what was used, not only in the US, but in much of Europe and Asia. GOSIPs main effect was in getting vendors to supply OSI.

GOSIP has also had a high political profile. Some folks believed that it would precipitate a shift away from TCP/IP to OSI. It never happened, and now OSI is moribund and GOSIP is changing in a drastic fashion.

GOSIP Version 1 was approved as a draft in October 1988, as a Full Use version in NIST Special Publication 500-187. It specified the OSI protocol stack. The current version is GOSIP Version 2 (FIPS 146-1; 3 April 1991).

GOSIP Version 2 was a watered down version. It contained the Connectionless Network Protocol, also known as ISO-IP. CLNP is the *Internet Protocol* (IP) with a bigger address space. But it isn't IP. Networks built from it won't interoperate with the Internet without conversion software. So GOSIP-2 required agencies to acquire products with two sets of protocols, only one of which, TCP/IP, would be used.

POSIT

On 16 September 1994, NIST entered a notice in the Federal Register for a FIPS 146-2 that renames GOSIP to *Profiles for Open Systems Internetworking Technologies* (POSIT).

This new FIPS removes the requirement to procure only OSI protocols. It does include a recommendation for what to use if an agency wants to acquire OSI, but it does not require acquiring OSI. For general use, POSIT refers to the Internet Standards specified by the *Internet Engineering Task Force* (IETF). The most specific reference is to RFC 1601, which describes the Internet Architecture Board (IAB), which is the parent organization of the IETF, and RFC 1602 which describes the Internet standardization process itself. These Internet Standards are of course the specifications of the TCP/IP protocols, including IP.

Translation: GOSIP is dead, and the U.S. Government is formally acknowledging that TCP/IP is preferable to OSI.

Bizarrely, ARPA was the principal funder for the original TCP/IP protocol suite. When Vint Cerf (now at MCI) and Bob Kahn (now at CNRI) came up with their ideas twenty years ago, ARPA, NASA, and NSF fostered the development of the protocols. The "big switch" on January 1, 1983, from the Network Control Protocol to TCP/IP was encouraged by the US agencies. And in 1983, the Defense Communications Agency bought in to the suite as well.

The phenomenal growth of the Internet is the direct result of the fact that it is built upon these protocols.

A historical view

A glance at the history explains why TCP/IP succeeded where OSI did not.

In 1977, the British Standards Institute proposed to ISO that a standard architecture was needed to define the communications infrastructure. (This, as with IFIP, CCITT and other efforts, shows how the road to hell is paved with good intentions. Because X.25 was unsatisfactory, the IFIP Working Group was set up in the hope that the technological community could forestall the political arena of ISO. It didn't.) ISO set up a subcommittee of a technical committee to study this [ISO/TC 97/SC 16].

The next year, 1978, ISO published its "Provisional Model of Open Systems Architecture" [ISO/TC 97/SC 16 N 34]. This was labelled a "Reference Model," and referred to as ISORM (ISORM—pronounced "eye-sorm" by Mike Padlipsky) [24]. In general, it was based on work done by Mike Canepa's group at Honeywell Information Systems, which came up with a seven-layered architecture, which itself owed much to IBM's proprietary *Systems Network Architecture* (SNA). SNA had been announced in 1974 and its seven layers don't correspond exactly to OSI/ISORM's. TC 97/SC 16 turned over proposal development to the American National Standards Institute (ANSI), to which Canepa and his technical lead, Charlie Bachman, presented their layered model. This, in turn, was the *only proposal* presented to the ISO subcommittee, at a meeting in Washington in March, 1978. It was accepted and published immediately.

A "refined" version of the ANSI submission to ISO appeared in June 1979. This published version is nearly identical to Honeywell's version of 1977. After an elaborate set of meetings, four International Standards were legislated. The extant NCP (host-host protocol) did not fit ISORM. ISO was trying to construct a nice set of geometric figures that would be a "tidy model." The original ARPANET crew was interested in getting things to actually work, to push bits around a system.

continued on next page

Protocol Wars: Is OSI Finally Dead? *(continued)*

John Quarterman remarked to me: "OSI specified before implementation. So specification took forever and implementation never really happened, except for bits and pieces. In addition, heavy government backing (by the EC—now the EU—and various national governments) led some OSI participants to attempt to substitute official authority for technical capability. OSI and IP started at about the same time (1977). OSI wandered off into the weeds and IP won the race. Those governments that backed OSI bet on the wrong horse."

Specification before implementation

In my opinion, this last is the answer: specification before implementation doesn't work. The necessity of delivering functioning hardware and software was what motivated both the ARPA team and the Network Working Group. This sort of group was not involved, beyond the very beginnings, in the bureaucracy of ISO and the creation of OSI.

The other problem was that TCP/IP was a suite with a real installed base. The implementations had evolved between 1974 and 1978 and were widely accepted within the technical community. By supporting a theoretical specification with no implementation, the PTTs of Europe and Japan were both sticking to the "old" telephony and telegraphy way of thinking at the same time they appeared to be mired in the "Not Invented Here" morass. The US government had funded the ARPANET and it was largely an American creation. A final reason that has been adduced is that the Japanese and European governmental representatives (all the PTTs were governmental) didn't want US manufacturers to have an unfair advantage, selling their extant TCP/IP products. Real money was involved; future profits were at stake.

The triumph of reality

The fact that local area networks and desktop workstations supported TCP/IP meant that there was an ever-increasing infrastructure built upon TCP/IP in Europe. The changing political situation in Eastern Europe and in the erstwhile USSR also meant increasing TCP/IP support (as all the specifications were in the public domain and thus no fees were payable). So far as I can tell, while OSI networking still exists in many places, it is not waxing; TCP/IP networks are.

The actual result was that many companies wasted a decade trying to produce OSI products while the networking community increasingly used TCP/IP. While Mischa Schwartz (in 1987) could still believe: "TCP and ISO TP class 4 will coexist for some time, but ... as more manufacturers and users begin to adopt the ISO suite of transport protocols, use of TCP will begin to decrease"; Fred Halsall, a staunch OSI supporter, has pointed out "In practice ... there are two major open system (vendor-independent) standards: the TCP/IP protocol suite and those based on the evolving ISO standards" [1992]. More manufacturers and users haven't adopted the OSI suite.

Now that NIST has succumbed, the war has been won by TCP/IP; I can only be aghast at the energy and time that went into the protocol wars, which have only ended late in 1994, with the US government's recognition of the de facto victory by TCP/IP.

References

[1] Jacobsen, Ole, "The Trouble with OSI," *ConneXions,* Volume 6, No. 5, May 1992.

[2] desJardins, Richard, "*Opinion:* OSI is (Still) a Good Idea," *ConneXions,* Volume 6, No. 6, June 1992.

[3] Metcalfe, Bob, Smart, Bob, and Blackshaw, Bob, "Letters to the Editor," *ConneXions,* Volume 6, No. 6, June 1992.

[4] Rose, Marshall, "Comments on: '*Opinion:* OSI is (Still) a Good Idea,'" *ConneXions,* Volume 6, No. 8, August 1992.

[5] desJardins, R., "Comments on: 'Comments on: *Opinion:* OSI is (Still) a Good Idea,'" *ConneXions,* Vol. 6, No. 10, October 1992.

[6] desJardins, Richard, "Internet 2000," *ConneXions,* Volume 6, No. 10, October 1992.

[7] Hoffmann, Harald, "Letter to the Editor," *ConneXions,* Volume 6, No. 11, November 1992.

[8] Hoffmann, Harald, "Letter to the Editor," *ConneXions,* Volume 7, No. 10, October 1993.

[9] Crocker, D. & Feit, S., "Letters to the Editor," *ConneXions,* Volume 7, No. 2, February 1993, p 23–25.

[10] Jacobsen, O., Karn, P., Auerbach, K., "Have you heard the GOSIP?" *ConneXions,*Volume 1, No. 1, May 1987.

[11] Huizer, Erik, "The IETF integrates OSI related work," *ConneXions,* Volume 7, No. 6, June 1993.

[12] *ConneXions,* Special Issue: IP—The Next Generation,Volume 8, No. 5, May 1994.

[13] Crocker, D., "Making Standards the IETF Way," *ConneXions,* Volume 7, No. 9, September 1993.

[14] Postel, J., ed., "Internet Protocol," RFC 791, September 1981.

[15] Vair, D., "Components of OSI: X.25—the Network, Data Link, and Physical Layers of the OSI Reference Model," *ConneXions,* Volume 4, No. 12, December 1990.

[16] Hagens, R., "Components of OSI: CLNP or A Day in the Life of Ivan CLNPacket," *ConneXions,* Volume 3, No. 10, October 1989.

[17] Huitema, C., "Charter of the Internet Architecture Board (IAB)," RFC 1601, March 1994.

[18] Internet Architecture Board and Internet Engineering Steering Group, "The Internet Standards Process—Revision 2," RFC 1602, March 1994.

[19] E. Huizer & D. Crocker, "IETF Working Group Guidelines and Procedures," RFC 1603, March 1994.

[20] Postel, J., "Internet Official Protocol Standards," RFC 1610, July 1994.

[21] Quarterman, Jon, "The History of the Internet and the Matrix," *ConneXions,* Volume 9, No. 4, April 1995.

[22] Stefferud, E., "A Plea For Internet Peace," *ConneXions,* Volume 1, No. 5, September 1987.

[23] Stefferud, E., "Paradigms Lost," *ConneXions,* Volume 8, No. 1, January 1994.

[24] Padlipsky, M., *The Elements of Networking Style,* Pretince-Hall 1985, ISBN 0-13-268111-0.

PETER H. SALUS is the author of *A Quarter Century of UNIX* (Addison-Wesley, 1994) and *Casting the Net: From ARPANET to Internet and Beyond* (Addison-Wesley, 1995). He holds B.Sc., M.A., and Ph.D. degrees from NYU. Salus is a member of the AAAS, a life member of the American Oriental Society, and the Internet Society, among other things.

The RA Route Server Service Overview

by Jessica Yu, Merit Network, Inc.

Introduction

This article describes the current *Route Server* (RS) service provided by the NSF-sponsored *Routing Arbiter* (RA) project. Route Servers are installed at each *Network Access Point* (NAP), the interconnection points where many *Internet Service Providers* (ISPs) and other networks exchange traffic. The purpose of the RS service at each NAP is to facilitate and simplify inter-domain routing among the various providers attached to the NAP. The RS service is not limited to the NAPs; it could be provided at any Internet interconnection point that shares the basic functionality of a NAP.

In this article, the terms NSP (Network Service Provider) and ISP (Internet Service Provider) are used interchangeably.

What is a NAP Route Server?

A NAP Route Server can be conceived as a system that performs routing processing for the ISPs connected to the NAP. The RS exchanges routing information with the routers attached to the NAP, and passes routing information from one ISP to another meeting their policy requirements, thus allowing direct traffic exchange between ISP routers at the NAP. The Route Server itself does not forward packets or perform any switching functions for the ISPs. (See Figure 1.)

The RA Route Server system is a Sun SPARC 20 workstation running SunOS. Special routing software developed by the RA project performs the routing exchange and processing functions described in detail below. The RS is operated by the RA project, a joint undertaking of Merit Network, Inc., and the University of Southern California Information Sciences Institute (ISI).

According to current plans, two RSs will serve each NAP to provide redundant service.

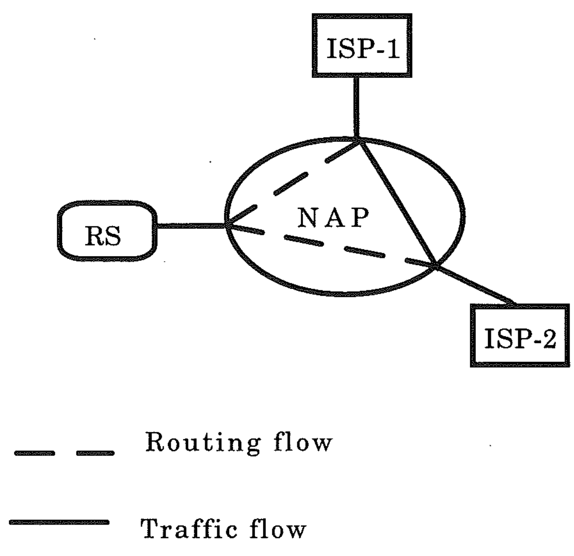

Figure 1: *Route Server facilitates information exchange while not involved in packet forwarding*

How does a Route Server work?

The RS facilitates routing exchange among the NAP-attached ISPs by gathering routing information from each ISP routers on the NAP, processing the information based on the ISP's routing policy requirements, and passing the processed routing information to each ISP router. Initially, the RS uses BGP-4 as the inter-domain routing protocol to exchange routing information with NAP-attached ISP routers.

As noted earlier, the RS does not forward packets among the NAP-attached ISPs. The RS uses BGP's third-party routing information capabilities to pass routing information from one ISP to another, with the next hop pointing to the ISP router that advertises the route to the RS. Traffic is therefore exchanged directly among the ISP routers on the NAP, even though the routing information is provided by the RS.

The RS has the ability to create a *Routing Information Base* (RIB), referred as a "View" in RS parlance, for each ISP router that peers with the RS. The view created for a given ISP maintains routing information which meets the policy requirements of that particular ISP. The view makes it possible for an ISP peering with the RS to obtain the same routing information from the RS that it would if it peered with every other ISP on the NAP, without the presence of the RS service. That is, the RS could give a different path towards a given destination to different ISPs, if such paths were available and if such policy were required by the ISPs. For example, ISP-1, ISP-2, ISP-A and ISP-B are all connected to a NAP with the RS. ISP-1 and ISP-2 can both reach destination X. ISP-A prefers to traverse ISP-1 to reach X, while ISP-B favors to traverse ISP-2 to reach destination X. The RS will provide routes to satisfy both ISP-A and ISP-B's policy needs. Similarly, if ISP-1 does not want to be the transit to destination X for the traffic coming from ISP-A, the RS will not pass ISP-1's route to ISP-A. (See Figure 2.)

Readers may refer to [1] for detailed information about the design of the Route Server.

In order for the RS to tailor its route processing to meet the policy requirements of an ISP, the ISP must register its inter-domain routing policy information in the RADB (*Routing Arbiter Database*) provided by the RA service. The RADB is a part of the IRR (*Internet Routing Registry*), a virtual database currently comprising databases provided by RA, RIPE, MCI, and CA*net. The RS will derive a given ISP's routing policy based on the information registered in the IRR. Please refer to [2] and [3] for information about the RADB.

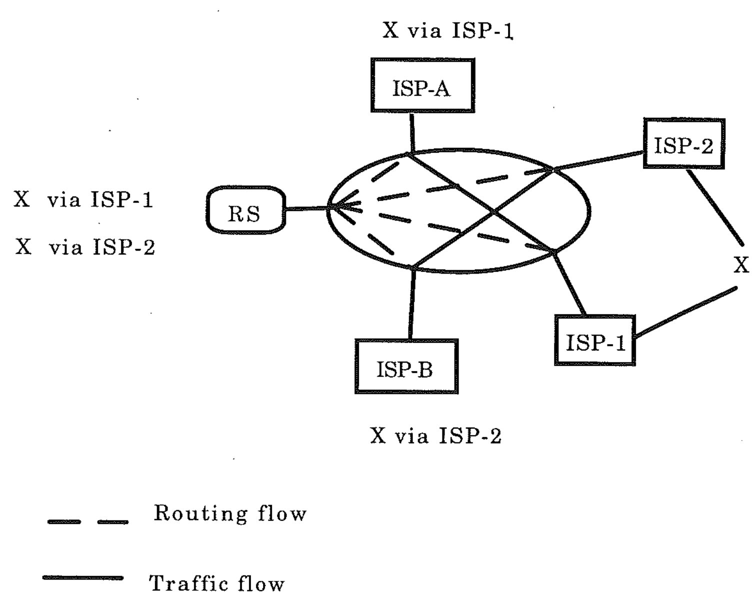

Figure 2: *Route Server does customized routing*

continued on next page

RA Route Server Service Overview *(continued)*

What does a RS offer?

• *Scalable Routing at NAPs:* The RS facilitates routing information exchange among the ISPs attached to a NAP by peering with each ISP router on the NAP. Instead of a full-mesh BGP peering among all the ISPs on the same NAP, the ISPs could peer only with the RS, and still achieve the goal of full routing exchange with the other ISPs. The RS thus reduces the number of peering sessions each ISP router needs to process from O(n) to O(1).

Note that in addition to peering with the RS, ISP routers may also optionally peer with each other if so desired.

• *Simplified Routing Processing on ISP Routers:* The RS processes routing information based on the routing policy described by each ISP. This includes, but is not limited to, route filtering for a particular ISP, selection of the desired path towards all destinations the ISP will be reaching, etc. This will greatly reduce routing processing, filtering and configuration management for ISP routers at the NAPs.

The RS will not distribute routes learned from one ISP to another ISP without the permission of both. This permission will be expressed in terms of the ISP's routing policy registered in the IRR.

• *Insertion of the RS Autonomous System (AS) Number in the Routing Path:* The RS can be configured to insert or suppress its own AS number in the AS path when passing routes from one ISP to another via BGP. This option is configurable on a peer-by-peer basis, and is configured according to the wishes of each ISP. This means that the Route Server could be viewed transparently when passing routing information.

• *Handling of Multi_Exit_Discriminator (MED):* The RS is able to pass along the Multi_Exit_Discriminator (MED) attribute defined in BGP protocol specification with no modification. Passing unmodified MED value allows traffic exchange between pair of ISPs using the MED as if they directly peered with each other.

The RS is also able to assign a MED to routes advertised by an ISP with a value specified by the ISP when passing them to other ISPs. This provides an ISP with the ability to apply desirable MED values towards selected ISPs.

• *Route Flapping Damping Mechanism:* A route flapping damping mechanism will be implemented in the RS to reduce the impact of frequent route flapping on ISP router performance.

• *Redundancy:* According to current plans, two RSs will serve each NAP to provide fault tolerance.

Why use the Route Server at a NAP?

Based on the RS functions described above, the RS at NAP offers the following:

• *Scalable Routing at NAPs:* As mentioned earlier, ISP routers on a NAP would need to establish full-mesh BGP peer sessions among themselves in order to exchange routing information without the presence of the RS. That is, if there were N ISPs present at a NAP, each would have N–1 peering sessions. When N is a large number, a sizable load could be placed on each router in order to maintain the required peering sessions and process the needed routing information. With the RS, each ISP only needs to peer with one RS—or two for backup purposes—instead of maintaining N–1 peer sessions.

• *Separation of Routing and Forwarding:* If a NAP did not have a Route Server, each ISP router would need to perform two major functions at all times: route processing and packet forwarding. A heavy traffic load could put a substantial extra burden, which would also need to process routing information. The load would be particularly heavy if the number of peering sessions was not small, the number of destination routes was large, and the policy was complicated. It would be ideal to have the routers concentrate on forwarding packets, and have another system handle routing. The Route Server achieves just this goal: it processes routing information for each ISP's router, thus enabling the ISP routers to concentrate on packet switching.

• *Simplified Routing Configuration Management on ISP's Routers:* The RS processes routing information based on the routing policy defined by each ISP. This includes route filtering, e.g., setting up routing firewalls, selecting the desired path towards all destinations the ISP will be reaching and other tasks. These routing tasks would normally be configured and implemented on the ISP routers. Therefore, the RS would greatly reduce the routing processing, filtering and configuration management load on the ISP routers at the NAPs.

It should be noted that RS not only can be used for some complicated routing policies but also can be used to facilitate simple routing policies. An ISP's policy could be as simple as to accept all the routes advertised by another ISP at a NAP.

• *Enforcing Good Routing Engineering:* The RS provides more flexibility in terms of adding new mechanisms or techniques to its routing code than many commercial vendors. Peering with the RS may therefore provide a quick fix to urgent problems. For example, route flapping consumes a great deal of precious router processing time, and is currently a major routing engineering concern. The Route Server can help reduce the effects of route flapping by implementing a route damping mechanism in the RS.

Who can use the RS Service?

All NAP-attached ISPs are entitled to RS services at this stage. The RS will peer with anyone who requests to peer with it, providing the ISP agrees with the conditions described in "The Routing Arbiter Peering Agreement." [4].

Technically, the ISP needs to meet certain conditions in order to peer with the RS. The ISP is also required to use the modern *Inter-Domain Routing Protocol* (IDRP) to exchange routing information with the RS and register its policy information in the IRR, so the RS can process routing information based on the particular ISP's routing policy.

How to peer with the Route Server

In order to peer with the RS, the ISP administrator should submit a "Route Server Peering Session Request Template" [5] via e-mail to `rs-peer@ra.net`. Upon receiving each request, the RS will be configured to peer with the router. The routing exchange policy will be based on the routing policy description of the AS associated with this router, and is expected to be registered in the IRR. In the absence of the related AS policy information in the IRR, the RS will be configured to peer with the router with a simple default policy. By default, the RS will not distribute the routes advertised by this peer to other ASs, nor will it distribute routes from other ISPs to the peer.

The requestor will need to make the proper configuration on its own NAP-attached router in order for its router to establish a peering session with the RS.

continued on next page

RA Route Server Service Overview *(continued)*

Summary

As mentioned above, the purpose of the NAP Route Servers is to facilitate and simplify routing among the NAP-attached service providers. Among the advantages of using the RS service at NAP are scalable routing and optimum use of an ISP router's CPU power cycle for packet switching, by leaving routing processing tasks to the RS. The RS function and service will be evolving with more operational experience and feedback from ISPs and the Internet community.

The RA Route Servers are currently deployed at MAE-East provided by MFS, the New York area NAP provided by Sprint, the Pac*Bell NAP in San Francisco, the Ameritech AADS NAP in Chicago, and at MAE-West in the San Francisco area.

Acknowledgements

Special thanks to Susan Harris of Merit for editing and polishing the text of this document. Many thanks to Jon Postel, Cengiz Alaettinoglu of ISI for their constructive and helpful comments.

References

[1] Chu, J.(Editor), "Route Server Design," `http://www.merit.edu/routing.arbiter/RA/route.server.html`, `ftp://www.isi.edu/pub/hpcc-papers/ra/rs-desgin-02feb95`

[2] "Routing Arbiter Database Overview," `http://www.merit.edu/routing.arbiter/RA/RADB.tools.docs.html` `ftp://ftp.ra.net/pub/radb/OVERVIEW`

[3] "How to Register in the RADB," `http://www.merit.edu/routing.arbiter/RA/RADB.tools.docs/register.html`

[4] "The Routing Arbiter Peering Agreement," `http://www.merit.edu/routing.arbiter/RA/rapa.html`

[5] "The RS Peering Request Template," `http://www.merit.edu/routing.arbiter/RA/.peering.template.html`

[6] D. Estrin, J. Postel, and Y. Rekhter, "Routing Arbiter Architecture," *ConneXions,* Volume 8, No. 8, August 1994.

[7] Yu, J., Chen, E. and Joncheray, L., "A Routing Design for the Initial ATM NAP Architecture," *ConneXions,* Volume 8, No. 11, November 1994.

[8] Rekhter, Y., and Li, T. (Editors), "A Border Gateway Protocol 4 (BGP-4)," RFC 1654, July 1994.

[9] Yu, J., Chen, E. and Joncheray, L., "Routing for the Initial ATM-NAPs In the Presence of the Special ANS/NSFNET Attachment," *ConneXions,* Volume 9, No. 1, January 1995.

[10] Adams, A., "The MERIT Policy-Routing Configuration System," *ConneXions,* Volume 7, No. 2, February 1993.

JESSICA (JIE YUN) YU received her B.S. in Computer Science from Shanghai Jiao Tong University in China and M.S. in Computer Engineering from The University of Michigan. She is currently a Senior Internet Engineer of Merit Network Inc. and is one of the co-PIs of the Routing Arbiter Award from NSF. She joined Merit in 1987 and was a key contributor in the initial engineering and deployment of the T1 and T3 NSFNET backbone and has since worked on the NSFNET engineering and operation. The major focus of her work has been NSFNET routing architecture design, engineering and implementation, and Internet advanced routing research. She has performed pioneering work on BGP and its deployment and is one of the architects of CIDR. She co-authored several RFCs and co-chairs the IETF CIDR Deployment Working Group. She is currently working on various Routing Arbiter related tasks including Routing Registry design, NAP routing design and ATM-NAP testing. She is also a member of FEPG (Federal Networking Engineering Planning Group). E-mail: `jyy@merit.edu`

Book Review

FDDI: A High Speed Network, by Amit Shah and G. Ramakrishnan. Prentice Hall, Englewood Cliffs, NJ 07632, 1994. ISBN 0-13-308388-8.

Relevant

In the midst of al the hype, frenzy, and fascination over ATM, one might be tempted to ask whether FDDI is still relevant. The simple truth is that FDDI not only remains relevant, but it is providing high speed networking solutions in scenarios where current generation ATM switching technology has not yet lived up to expectations. Certain vBNS NAPs, commercial interconnects, and enterprise internet equivalents use FDDI to interconnect high-end routers in those circumstances where the number of T3 circuits terminated at a single site exceeds any single router's capacity. FDDI is also used as a campus backbone solution and as an upgrade LAN technology for NOS servers in enterprise nets. And FDDI is certain to play a role in the *evolution to switched internetworking.* So I'd say FDDI is relevant and will remain so for some time.

The criteria for whether a book on FDDI remains relevant are somewhat different than those on ATM. Books covering ATM may still be relevant if they focus on standards and theory, but FDDI standards are nearly 10 years old, and papers describing the Time Token Rotation protocol on which FDDI is based appeared in 1982, so one would hope that any book competing for shelf space today would address the practical aspects of FDDI.

Organization

Thankfully, Amit Shah and G. Ramakrishnan do a rather good job providing relevant material in this book. In fact, you have to look hard to find in-line references to standards. Four chapters describe FDDI nodes and topologies, the MAC, physical, and physical media dependent layers in ample if not painstaking detail, and diagrams modeling operational flows complement the text nicely. The chapter on *Station ManagemenT* (SMT) does an admirable job of explaining link and node management in as close to "plain-speak" as one could expect. The practical aspects of FDDI—issues to consider when selecting media, installing cabling, and when choosing a topology—are covered in the closing chapters. The final chapter provides guidelines for troubleshooting FDDI networks. These are probably the most interesting and valuable chapters, especially for those who are installing FDDI.

The typical *ConneXions* reader may find the chapter describing internetworking with FDDI rather mundane, and may not appreciate the somewhat lackadaisical treatment of SNMP and CMIP in the chapter on remote network management (the authors show a bias towards SMT which is not universally shared, and relegate SNMP, et. al. to a subsection entitled "FDDI Management with Other Protocols"). The propensity of the authors to intermingle discussion of Internet and OSI technology without clearly distinguishing between the two was something of a frustration, and the introductory chapters are disconnected, but I would not denounce this book on the basis of "occasional inaccuracies" and a slow start.

Good on balance

In the Preface, the authors promise "a book which is neither too technical nor too simplistic." On balance, Shah and Ramakrishnan deliver. If you have no knowledge or experience with FDDI, you should consider adding this book to your library.

—David M. Piscitello,
Core Competence, Inc.
dave@corecom.com

Announcement and Call for Papers

The *7th Joint European Networking Conference* (JENC7) will be held in Budapest, May 13–16 1996. The event is organized by TERENA, the Trans-European Research and Education Networking Association, with the local assistance of the Hungarian Academy of Sciences and HUNGARNET. The theme of the conference is "The Role of Research Networking in the Information Society." In line with the established tradition of the previous JENCs this 7th JENC aims at bringing together individuals from research and education, industry and government who are involved with planning, developing, implementing, managing, funding, and using national, regional and international computer networks for a four day meeting in which state of the art networking issues will be presented, discussed and demonstrated.

Topics

The following subject areas define, not exclusively, possible topics for paper submission:

- *User Support and Education:*
  Support of tele-collaboration
  Publishing issues in the Information Society
  Globalization of user support services
  Virtual education and learning communities
  Networked scientific research and its applications
  K–12
  Work and play in Cyberspace
  User education tools
  Networked information retrieval
  Library access in the Information Society

- *Policy, Economic and Societal Issues:*
  Networking developments in technologically emerging countries
  Technology transfer to technologically emerging countries
  Funding and pricing models for networks and networking services
  Commercialization, privatization and public access
  Electronic communities and the law
  Copyright and intellectual property rights issues
  Privacy and data protection
  Governments and the Information Society
  Effects of Telecommunication Liberalization

- *Network Engineering:*
  Building the Information Society
  Application of network technology to provide networking services
  International interoperability and network management issues
  Network management systems and methods
  Reliability, performance and scaling issues
  Network security mechanisms and incident handling

- *Network Technology:*
  New and international open network protocols
  Transmission, routing, and transport technologies
  Multicast developments
  High speed WANs
  Mobility developments

- *Application Technology:*
  Network application infrastructure
  Computer supported collaborative work
  Interoperability of application services
  Distributed applications' management
  Security aspects of distributed applications

- *Infrastructure developments:*
  European backbone developments
  ATM
  34 Mbps and beyond

Submissions

All papers must be written in English. Electronic submission is highly recommended and should take place as follows:

- ASCII or uuencoded *PostScript:* send by e-mail to: `jenc7-submit@terena.nl`
- *PostScript* documents: send by anonymous FTP to Internet host `erasmus.terena.nl` (IP address 192.87.30.2), into the directory `pub/jenc7/submit`. Please note that files deposited in this directory can only be written once and cannot be deleted afterwards.

There will be the opportunity for participants to present exciting applications of networking services in the form of a demonstration. Proposed demonstrations should be documented with a description not exceeding one page.

Important dates

Full manuscript due:	November 19, 1995
Proposals for demonstrations due:	November 19, 1995
Notification of acceptance to authors:	January 15, 1996
Camera-ready papers due:	March 31, 1996

Publications

Conference proceedings containing full papers will handed out to the participants. A selection of the best papers will be published as a special issue of *Computer Networks and ISDN Systems.*

Workshop and tutorials

The traditional *Network Technology Training Workshop* will be held the week prior to the conference. Travel and tuition support may be available for selected attendees. Additional information is available from the JENC7 Secretariat at the address below. Tutorials are planned to be held on May 13 (morning) and May 16 (afternoon) as well as May 17 (full day).

Working group presentations

There will be provisions for presentations describing the activities in TERENA Working Groups and IETF Areas. The participants in the EC's 4th FW Programme are invited to present progress reports and/or planned activities—in the realm of the TERENA coordinated SCIMITAR project.

Exhibition

An exhibition area will be available for International and Hungarian companies and institutions for demonstration of their products and services.

Venue

The premises of the Hungarian Academy of Sciences, one of the most beautiful historical buildings of Budapest, will accommodate the plenary sessions, most parallel technical sessions and the networking facilities of the conference. The TERENA Working Groups will also meet here. Within 3 minutes walk of the Academy of Sciences, the Danube Palace provides several attractive meeting rooms in neobaroque style. This building will house the rest of the technical sessions, special meetings, the demonstrations and lunch.

More information

JENC7 Secretariat
c/o TERENA Secretariat
Singel 466–468
NL-1017 AW Amsterdam
THE NETHERLANDS

`jenc7-sec@terena.nl`

`http://www.terena.nl/terena/jenc7/`

JENC7 Local Organization
c/o MTA SZTAKI
Kende u. 13–17
H-1111 Budapest
HUNGARY

`richter@sztaki.hu`

`http://www.iif.hu/jenc/`

Internet Survey Reaches 6.6 Million Host Level

First Half 1995 Growth is 37 Percent

Reston VA, USA, August 2, 1995—The latest results from the Internet's most basic and longest continuing measurement of its size were just released by Mark Lottor of Network Wizards of Menlo Park, CA, USA. The *Domain Survey* attempts to discover every host on the Internet by doing a complete search of the Domain Name System. The latest results were gathered during late July 1995. The data is available in the `zone` directory on `ftp.nw.com`, or `http://www.nw.com.`

50,000 networks

The Internet is a very complex, dynamic, distributed aggregation of more than 50 thousand autonomous networks. It defies definitive measurement. Nonetheless, these values constitute prima facie connected host computers, even if many might not always be reachable, and Lottor has carefully conducted these counts over many years. They are also extremely valuable for relative comparisons.

The Internet Research Task Force (IRTF) *Survey Working Group* (SWG) is analyzing these statistics and methodologies to detect anomalies and produce additional information.

Highlights

Newsworthy, strategic highlights of the most current values include:

- Very slightly decreased, but continued strong exponential growth rate. At the average rate of increase over the past 14 quarters, the total projected hosts at the end of the decade is 101 million.
- Hosts in 106 country domains were counted, an increase in 15 countries. (Note that verification has not been performed to verify that these hosts are physically located in the country.)
- The global commercial domain `.COM` continues not only to be the largest, but continues growing at a rapid rate.
- Germany and Japan are exhibiting very rapid growth rates among industrialized countries with a first half rates of 41% and 40%, respectively.
- In absolute terms, the USA had the largest jump of about 1,090,000 hosts—a rate of 24%. The USA values are subject to inherent uncertainties because of the mix of 3-letter global domains and the `.US` domain.
- Strong Russian Federation growth continues at a 68% half year rate.
- Most regional growth rates throughout the world continue at averages exceeding 40 percent.

Details

Updated color graphs of these trends, including those for most countries are available at:

`ftp://ftp.isoc.org/isoc/charts/hosts4.ppt` (PowerPoint v.4)
`ftp://ftp.isoc.org/isoc/charts/hosts3.ppt` (PowerPoint v.3)

About ISOC

The Internet Society is an International individual membership organization for the Internet global cooperation. Its International Secretariat can be reached at:

12020 Sunrise Valley Drive, Suite 210
Reston, VA, USA
`isoc@isoc.org` `http://www.isoc.org`
Tel: +1 703 648 9888 Fax: +1 703 648 9887

—A.M. Rutkowski, Executive Director, ISOC

Call for Papers

The NetWorld+Interop US Program Committee is pleased to solicit original technical papers for the 3rd annual *Interop Engineers' Conference,* held in conjunction with the NetWorld+Interop Conference and Exhibition, from April 1st through 5th, 1996.

Topics

In order to focus discussion and interaction, this year the Engineers' Conference is focusing on six topic areas of interest in computer-communications:

- Resource Management over Heterogeneous Networks
- Cell-based Routing
- Traffic management and the future of Congestion Control
- Distributed Applications Management
- Video over Enterprise Networks
- High-speed Packet Filtering and Firewalling

Submissions

A detailed description of each topic area can be found at URL `http://www.interop.com`. This conference seeks to bring together research scholars, engineers, and vendors to address pragmatic engineering issues in the field of networking and distributed systems interoperability. It is an excellent forum for engineers and researchers to publish papers on solutions to today's engineering-related problems.

Interested parties should submit abstracts of their papers by *September 8, 1995.* An abstract should be 500–1,000 words in length and convey the key aspects of the paper. All abstracts should be submitted in ASCII. The program committee will indicate its acceptance (or not), no later than *September 22, 1995.* To submit an abstract, send a message:

```
To: engrconf@interop.com
Subject: abstract
```

(Do not put anything else in the Subject: line.) The message should contain your complete contact information (name, affiliation, postal address, telephone, facsimile, and e-mail) along with your abstract. An automated reply will confirm receipt of your abstract.

If an abstract is accepted, the author(s) should submit a first draft of their paper by *December 31, 1995.* A paper should be between 10 to 16 pages in length, and be written in technical English. All papers should be submitted either in ASCII or *PostScript.* The program committee will indicate its acceptance (with comments) or not, on *January 19, 1996.*

Publication and presentation

If a paper is accepted, the author(s) should submit the final copy of their paper, reflecting the comments of the program committee by *February 23, 1996.* All final copies will be published in the event proceedings. Upon receipt of the final copy, the program committee will inform the author(s) if their papers are to be presented at the event. Presentation should be 20–25 minutes, excluding questions.

Note that although every author who submits a final copy of an accepted paper receives a complimentary admission to the Engineers' Conference as well as the NetWorld+Interop General Conference and Exhibition, there may not be sufficient speaking slots for each accepted paper.

NetWorld+Interop expands to London and Sydney

SOFTBANK Exposition and Conference Company recently announced that they will be adding London and Sydney to their conference and exhibition "World Tour." This brings to seven the total number of NetWorld+Interop events for 1996. On October 28 through November 1, 1996 networking and interoperability hardware, software and transport technology products, services and applications will converge at Earls Court 2 for *NetWorld+Interop 96 London,* the first UK-based NetWorld+Interop Conference and Exposition. Down under, a similar event entitled *NetWorld+Interop 96 Sydney* will be held November 25–29, 1996 at the Sydney Exhibition and Conference Centre, Darling Harbour.

Hyperspeed

"The hyperspeed of the development of this market requires an ongoing, global arena in which network professionals can learn, touch and test the latest products and technologies. NetWorld+Interop attendees will be able to talk directly with technology experts and executives from leading companies and gain valuable hands-on experience with the newest products on our interactive show floor," explained Michael D. Millikin, Sr. Vice President, NetWorld+Interop. "These are also the ideal venues for vendors who need to reach UK-based and Australia-based network computing professionals."

Program

A full schedule of tutorials, conferences, workshops, technology demonstrations, and exhibits will provide attendees with access to the largest pool of networking hardware, software and transport products, services and applications assembled under one roof—at all seven events.

NetWorld+Interop is the most comprehensive event in the networking industry. It is the summit for interoperability, bringing together the products, services and applications networking professionals need and offering the highest quality courses led by industry experts, including the original Internet developers, technology inventors and leading reference authors. This event was firmly established as "The Networking Summit" with its 1994 debut, and is the global gathering place for the industry's best and brightest networking professionals. The NetWorld+Interop 1994 World Tour was a success with events in Las Vegas, Berlin, Tokyo, Atlanta and Paris.

Partners

SOFTBANK Expos has retained Real Time Events Ltd to jointly produce NetWorld+Interop 96 London. Real Time Events has over 50 years of combined trade show experience with its IT event professionals. In Australia NetWorld+Interop 96 Sydney will be produced jointly with Synergy Conventions. Synergy Conventions produces more than 100 conferences in Australia, including information technology and telecommunications conferences.

For information on exhibiting at NetWorld+Interop 96 London call David Conn at Real Time Events on +44 181 849 6260. For information on exhibiting at NetWorld+Interop 96 Sydney call Elena Cohen at Synergy Conventions on +61 2 369-1242.

NetWorld+Interop 96 Sydney is the only Australian networking event, following the discontinuation of IDG's Network World in 1995. IDG will become an official co-sponsor of the NetWorld+Interop Sydney event and additional co-sponsorships are being defined.

More information

NetWorld+Interop information is available on-line through the World-Wide Web at `http://www.interop.com`.

Future NetWorld+Interop Dates and Locations

NETWORLD+INTEROP 95 N+I

NetWorld+Interop 95	Paris, France	September 11–15, 1995
NetWorld+Interop 95	Atlanta, GA	September 25–29, 1995
NetWorld+Interop 96	Las Vegas, NV	April 1–5, 1996
NetWorld+Interop 96	Frankfurt, Germany	June 10–14, 1996
NetWorld+Interop 96	Tokyo, Japan	July 15–19, 1996
NetWorld+Interop 96	Atlanta, GA	September 16–20, 1996
NetWorld+Interop 96	Paris, France	September 23–27, 1996
NetWorld+Interop 96	London, England	Oct 28–Nov 1, 1996
NetWorld+Interop 96	Sydney, Australia	November 25–29, 1996

All dates are subject to change.

More information

Call 1-800-INTEROP or 1-415-578-6900 for more information. Or send e-mail to `info@interop.com` or fax to 1-415-525-0194. For the latest information about NetWorld+Interop including *N+I Online!* as well as other SOFTBANK produced events, check our home page at `http://www.interop.com`

NetWorld+Interop is produced by SOFTBANK Exposition and Conference Company, 303 Vintage Park Drive, Foster City, California 94404–1138, USA.

Write to *ConneXions*!

We'd love to hear your comments, suggestions and questions about anything you read in *ConneXions*. Our editorial address is given below. Use it for letters to the Editor, requests for the index of back issues, questions about particular articles etc.:

ConneXions—The Interoperability Report
303 Vintage Park Drive
Suite 201
Foster City
California 94404–1138
USA
Phone: +1 415-578-6900 or 1-800-INTEROP (Toll-free in the USA)
Fax: +1 415-525-0194
E-mail: `connexions@interop.com`
URL: `http://www.interop.com`

Subscription information

For questions about your subscription please call our customer service hotline: 1-800-575-5717 or +1 610-565-6864 outside the USA. This is the number for our new subscription agency, Seybold Publications. Their fax number is +1 610-565-1858. The mailing address for subscription payments is: P.O. Box 976, Media, PA 19063–0976.

CONNEXIONS
303 Vintage Park Drive
Suite 201
Foster City, CA 94404-1138
Phone: 415-578-6900
FAX: 415-525-0194

ADDRESS CORRECTION
REQUESTED

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1

2094C:9002083 12/99
BREWSTER KAHLE
WAIS INC
690 5TH ST
SAN FRANCISCO CA 94107-1517

Printed on recycled paper